IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
TYLER DIVISION

REALTIME DATA LLC d/b/a IXO, Plaintiff, v. DELL INC., EMC CORPORATION, ILAND INTERNET SOLUTIONS CORPORATION, AND VEEAM SOFTWARE CORPORATION Defendants.	Case No. 6:16-cv-89 JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT AGAINST DELL INC., EMC CORPORATION, ILAND INTERNET SOLUTIONS CORPORATION, AND VEEAM SOFTWARE CORPORATION

This is an action for patent infringement arising under the Patent Laws of the United States of America, 35 U.S.C. § 1 *et seq.* in which Plaintiff Realtime Data LLC d/b/a IXO ("Plaintiff," "Realtime," or "IXO") makes the following allegations against Defendants Dell Inc. ("Dell"), EMC Corporation ("EMC"), Iland Internet Solutions Corporation ("Iland"), and Veeam Software Corporation ("Veeam") (collectively, "Defendants"):

PARTIES

1. Realtime is a New York limited liability company. Realtime has places of business at 5851 Legacy Circle, Plano, Texas 75024, 1828 E.S.E. Loop 323, Tyler, Texas 75701, and 116 Croton Lake Road, Katonah, New York, 10536, and is organized under the laws of the State of New York. Realtime has been registered to do business in Texas since May 2011. Since the 1990s, Realtime has researched and developed specific solutions for data compression, including, for example, those that increase the speeds at

which data can be stored and accessed. As recognition of its innovations rooted in this technological field, Realtime holds over 45 United States patents and has numerous pending patent applications. Realtime has licensed patents in this portfolio to many of the world's leading technology companies. The patents-in-suit relate to Realtime's development of advanced systems and methods for fast and efficient data compression using numerous innovative compression techniques based on, for example, particular attributes of the data.

2. On information and belief, Defendant Dell is a Delaware corporation, with its principal place of business at One Dell Way, Round Rock, Texas 78682. On information and belief, Dell has a large services and data center location in Plano, Texas.[1] On information and belief, Dell can be served through its registered agent, Corporation Service Company, 211 East Seventh Street, Suite 620, Austin, Texas 78701-3218.

3. On information and belief, Defendant EMC is a Massachusetts corporation, with its principal place of business at 176 South Street Hopkinton, MA 01748. On information and belief, EMC can be served through its registered agent, CT Corporation System, 1999 Bryan St., Ste. 900, Dallas, TX 75201. On information and belief, EMC has known about Realtime's patent portfolio, years before it starting infringing Realtime's patents. Specifically, in February of 2004, EMC received a letter regarding Realtime's innovative digital-data compression technology. That letter also attached 5 issued U.S. Realtime Patents, which taught various embodiments of digital-data compression. On information and belief, EMC reviewed that letter and the attachments during the following 30 days, and sent a responsive letter to Realtime on March 15, 2004, stating that it refused to discuss licensing Realtime's patented technology. Years later, EMC began to infringe Realtime's patented technology, in numerous ways, including

[1] http://www.dell.com/content/topics/global.aspx/sitelets/solutions/perot/contact_us?c=us&l=en&cs=RC966726

those discussed in this Complaint.

4. On information and belief, Defendant Iland is a Texas corporation, with its principal place of business at 1235 North Loop West, Suite 800, Houston, TX 77008.[2] On information and belief, Defendant Iland operates cloud data centers at 1649 West Frankford Rd., Carrollton, Texas[3] and 12001 North I-45, Houston, Texas.[4] On information and belief, Iland can be served through its registered agent, Brian Ussher, 1235 North Loop West, Suite 800, Houston, TX 77008.

5. On information and belief, Defendant Veeam is a Delaware corporation, with its principal place of business at 2520 Northwinds Parkway, Suite 600, Alpharetta, GA 30009. On information and belief, Veeam can be served through its registered agent, CT Corporation System, 1300 East Ninth Street, Cleveland, Ohio 44114.

6. On information and belief, Defendant Dell has a business alliance with Defendant EMC that includes customization of Dell products to work together with EMC products such as EMC Data Domain.[5] Dell and EMC also market joint products such as the Dell / EMC DD Series deduplication storage systems (including DD140, DD610, DD630, and DD670), which are appliances incorporating EMC's Data Domain technology, including Global Compression™ together with Stream Informed Segment Layout (SISL™) Scaling Architecture.[6] On information and belief, Defendant Dell will complete its acquisition of Defendant EMC later this year.[7] Defendant Dell thus promotes the use of EMC products together with Dell's own products and services. On information and belief, these arrangements between EMC and Dell are based on ongoing

[2] http://www.iland.com/contact/
[3] http://www.iland.com/services/locations/iland-cloud-datacenter-dallas-texas/
[4] http://www.iland.com/services/locations/iland-cloud-datacenter-houston-texas/
[5] https://www.emc.com/data-protection/data-domain/data-domain-operating-system/dell-integrations.htm
[6] http://www.dell.com/downloads/global/products/pvaul/en/dell-emc-dd-series-brochure.pdf
[7] http://www.usatoday.com/story/tech/2015/10/12/dell-buy-emc-largest-tech-deal-ever/73727530/

contractual agreements between them. As further explained below, EMC Data Domain and related products such as Dell / EMC DD Series deduplication storage systems (including DD140, DD610, DD630, and DD670) infringe the asserted patents. Accordingly, Dell and EMC are properly joined in this action pursuant to 35 U.S.C. § 299.

7. On information and belief, Defendant Iland has a business alliance with Defendant Veeam that includes customization of Iland products to work together with Veeam products, such as iland Cloud Backup using Veeam Cloud Connect.[8] Defendant Iland thus promotes the use of Veeam products together with Iland's own products and services. On information and belief, these arrangements between Iland and Veeam are based on ongoing contractual agreements between them, pursuant to which Iland acts as a Veeam Cloud & Service Provider (VCSP) partner.[9] As further explained below, Veeam Backup & Replication (which includes Veeam Cloud Connect) infringe the asserted patents. Accordingly, Iland and Veeam are properly joined in this action pursuant to 35 U.S.C. § 299.

8. On information and belief, Defendant Veeam has a business alliance with Defendant EMC, pursuant to which Defendant Veeam offers a "Storage Availability Solution" that integrates Veeam products such as Veeam® Availability Suite™ with EMC storage products such as VNX/VNX2/VNXe storage systems.[10] As further

[8] http://www.iland.com/services/iland-cloud-backup-with-veeam/

[9] https://www.veeam.com/wp-cloud-connect-reference-architecture-veeam-backup-replication-v8.html ("One of the new and greatest features of Veeam® Backup & Replication™ v8 is Veeam Cloud Connect. With it, Veeam users can easily send backup copies offsite to remote locations managed by Veeam Service Providers. Also, Service Providers can use Cloud Connect to build their own remote repositories and offer their customers Backup Storage as a Service."); https://www.veeam.com/find-a-veeam-cloud-provider.html (identifying Iland as one of many "Veeam Cloud & Service Provider partners with datacenters in United States").

[10] https://www.veeam.com/emc-storage-availability-solution.html ("Veeam® Availability Suite™ combines with EMC storage to provide integration with VNX/VNX2/VNXe primary storage snapshots and support for Data Domain Boost; increasing data

explained below, Veeam® Availability Suite™, which includes Veeam Backup & Replication™,[11] as well as EMC VNX2 Series products, infringe the asserted patents. Accordingly, EMC and Veeam are properly joined in this action pursuant to 35 U.S.C. § 299.

JURISDICTION AND VENUE

9. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has original subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

10. This Court has personal jurisdiction over Defendant Dell in this action because Dell has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Dell would not offend traditional notions of fair play and substantial justice. Dell, directly and through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents. Dell is registered to do business in the State of Texas and has appointed Corporation Service Company, 211 East Seventh Street, Suite 620, Austin, Texas 78701-3218 as its agent for service of process. In addition, Dell has a principal place of business and a large services and data center location in Texas.

11. This Court has personal jurisdiction over Defendant EMC in this action

---

availability, enhancing secondary operation performance and opening new storage solution capabilities. **Veeam improves availability with EMC® storage solutions integration:** Veeam's integration with EMC VNX, VNX2 and VNXe hybrid storage arrays adds intelligence to EMC storage snapshots, increases backup speed, allows granular recovery of virtual machines (VMs) and application items, and enables patch and configuration testing.").

[11] https://www.veeam.com/availability-suite-licensing-faq.html ("Veeam® Availability Suite™ combines the industry leading backup, restore and replication capabilities of Veeam Backup & Replication™ with the advanced monitoring, reporting and capacity planning functionality of Veeam ONE™.")

because EMC has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over EMC would not offend traditional notions of fair play and substantial justice. EMC, directly and through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents. EMC is registered to do business in the State of Texas and has appointed CT Corporation System, 1999 Bryan St., Ste. 900, Dallas, TX 75201 as its agent for service of process. In addition, EMC has offices in Austin, Dallas, and Houston in Texas.[12]

12. This Court has personal jurisdiction over Defendant Iland in this action because Iland has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Iland would not offend traditional notions of fair play and substantial justice. Iland, directly and through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents. Iland is registered to do business in the State of Texas and has appointed Brian Ussher, 1235 North Loop West, Suite 800, Houston, TX 77008 as its agent for service of process. In addition, Iland has a principal place of business in Houston, Texas and operates cloud data centers in Carrollton, Texas and Houston, Texas.

13. This Court has personal jurisdiction over Defendant Veeam in this action because Veeam has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Veeam would not offend traditional notions of fair play and substantial justice. Veeam, directly and through subsidiaries or intermediaries, has committed and

[12] https://jobs.emc.com/location/united-states-texas-jobs/414/6252001-4736286/3

continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents. Veeam is registered to do business in the State of Texas.

14. Venue is proper in this district under 28 U.S.C. §§ 1391(b), 1391(c) and 1400(b). Each of Defendants Dell, EMC, Iland, and Veeam is registered to do business in Texas, and upon information and belief, has transacted business in the Eastern District of Texas and has committed acts of direct and indirect infringement in the Eastern District of Texas. In addition, Defendant Dell has a principal place of business in Texas, Defendant EMC has places of business in Texas, and Defendant Iland has a principal place of business and cloud data centers in Texas.

COUNT I

INFRINGEMENT OF U.S. PATENT NO. 7,161,506

15. Plaintiff realleges and incorporates by reference paragraphs 1-14 above, as if fully set forth herein.

16. Plaintiff Realtime is the owner by assignment of United States Patent No. 7,161,506 ("the '506 patent") entitled "Systems and methods for data compression such as content dependent data compression." The '506 patent was duly and legally issued by the United States Patent and Trademark Office on January 9, 2007. A true and correct copy of the '506 patent, including its reexamination certificates, is included as Exhibit A.

Dell Rapid Recovery Software

17. On information and belief, Dell has, or will soon have, made, used, offered for sale, sold and/or imported into the United States Dell products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Dell's compression products and services, such as, *e.g.,* the Rapid Recovery software product (v6.0.1),[13] the AppAssure v5.4 software

[13] http://software.dell.com/landing/6178

product,[14] the Dell DL1000 Backup and Recovery Appliance,[15] the Dell PowerVault DL4000 Backup and Recovery Appliance,[16] the Dell DL4300 Backup and Recovery Appliance,[17] the Dell DR4100 Disk Backup Appliance, the Dell DR6000 Disk Backup Appliance,[18] the Dell DR2000v backup disk virtual appliance,[19] the Dell SonicWALL WAN Acceleration Virtual Appliance (WXA) 5000,[20] and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

18. On information and belief, Dell has directly infringed and continues to infringe the '506 patent, for example, through its own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the '506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Dell uses the Accused Instrumentality to practice infringing methods for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing

[14] http://documents.software.dell.com/appassure/5.4.1/user-guide/introduction-to-appassure-5/about-appassure-5
[15] http://software.dell.com/products/appassure-dl1000-backup-and-recovery-appliance/
[16] https://partnerdirect.dell.com/sites/channel/Documents/Dell-PowerVault-DL4000-Spec-Sheet.pdf
[17] http://software.dell.com/products/dl4300-backup-and-recovery-appliance/
[18] http://software.dell.com/products/dr-series-disk-backup-appliances/
[19] http://software.dell.com/products/dr2000v-virtual-backup-appliance/
[20] http://accessories.dell.com/sna/PopupProductDetail.aspx?c=us&l=en&cs=04&sku=A7004425&price=4,495.00&client=config

technical support and repair services for the Accused Instrumentality to Dell's customers.

19. The Accused Instrumentality is a computer-implemented method for data compression. This system minimizes the amount of data transmitted over a network and stored on a backup device. The Accused Instrumentality employs several data compression techniques to achieve this goal.

20. The Accused Instrumentality analyzes data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types:

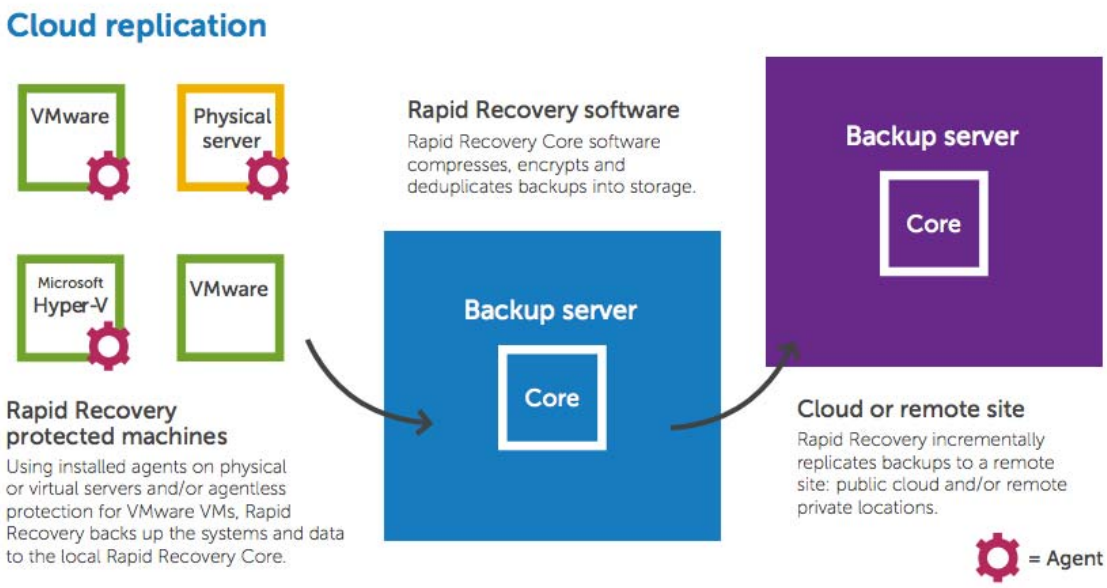

See http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf. See also http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf ("The new Rapid Recovery Repository (R3) is powered by the same leading-edge technology as Dell DR series backup and deduplication appliances."); http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf:

> The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

21. The Accused Instrumentality performs content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified. *See, e.g.*, http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf ("The R3, powered by the best-of-breed data deduplication technology platform used in Dell DR series appliances, can reside on separate Windows or Linux servers … Client-side data deduplication (also called source-side deduplication) uses the agent to remove redundant backup data before transmitting data to the R3. Using client-side data deduplication in tandem with target deduplication greatly reduces the amount of data sent over a LAN."); http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("AppAssure becomes the Rapid Recovery product under a Dell Data Protection (DDP) brand. The product is block- and snapshot-based, only storing unique, changed blocks with a five-minute RPO."). See also http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf:

> The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

22. The Accused Instrumentality performs data compression with a single data compression encoder, if a data type of the data block is not identified. *See, e.g.*,

http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("Rapid Recovery feature list: … There is a new deduplication and compression engine from the DR target-based product set, meaning Dell now has a single dedupe engine across its portfolio."); http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf ("The Dell DR4000 has made further optimizations, and can actually dedupe and compress as part of the same inline process. This provides the benefits of compression without requiring that space be dedicated to staging uncompressed data.").

23. In the Accused Instrumentality analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. See also http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf:

> The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

24. On information and belief, Dell also directly infringes and continues to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

25. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way. In particular, similar deduplication and compression technology used in Dell's DR backup appliances (e.g. DR4100, DR6000, and DR2000v) is used in Dell's Rapid Recovery software, Dell's AppAssure software, and in Dell's appliances built using Dell's AppAssure software (e.g. the Dell DL1000 Backup and Recovery Appliance, the Dell PowerVault DL4000 Backup and

Recovery Appliance, and the Dell DL4300 Backup and Recovery Appliance). See, e.g., http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("Rapid Recovery feature list: … There is a new deduplication and compression engine from the DR target-based product set, meaning Dell now has a single dedupe engine across its portfolio."); http://software.dell.com/products/appassure-dl1000-backup-and-recovery-appliance/ ("Built by Dell and powered by AppAssure, the DL1000 provides fast backup and restores on virtual machines, physical servers and in the cloud to enable local, offsite or disaster recovery."); http://software.dell.com/documents/dell-appassure-replication-technicalbrief-29890.pdf ("Dell™ AppAssure™ delivers advanced, flexible replication options to protect any organization, along with compression, deduplication and encryption"); http://documents.software.dell.com/appassure/5.4.3/user-guide-revision-b-english/configuring-the-appassure-core/managing-a-repository/modifying-repository-settings ("Enable Deduplication: Clear this checkbox to turn off deduplication, or select this checkbox to enable deduplication; Enable Compression: Clear this checkbox to turn off compression, or select this checkbox to enable compression.").

26. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

27. On information and belief, Dell has had knowledge of the '506 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, Dell knew of the '506 patent and knew of its infringement, including by way of this lawsuit.

28. Upon information and belief, Dell's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary

way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Dell instructs users of its Rapid Recovery software that "Rapid Recovery Core software compresses, encrypts and deduplicates backups into storage." See http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf. For similar reasons, Dell also induces its customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. Dell specifically intended and was aware that these normal and customary activities would infringe the '506 patent. Dell performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Dell engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Dell has induced and continue to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '506 patent, knowing that such use constitutes infringement of the '506 patent.

29. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Dell has injured Realtime and is liable to Realtime for infringement of the '506 patent pursuant to 35 U.S.C. § 271.

30. As a result Dell's infringement of the '506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Dell's infringement, but in no event less than a reasonable royalty for the use made of the invention by Dell, together with interest and costs as fixed by the Court.

EMC Data Domain

31. On information and belief, EMC has made, used, offered for sale, sold and/or imported into the United States EMC products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, EMC's compression products and services, such as, *e.g.*, the EMC Data Domain product, the EMC VNX2 Series products (including the EMC VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 products),[21] the EMC XtremIO Storage Array,[22] and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

32. On information and belief, Dell has made, used, offered for sale, sold and/or imported into the United States combination products incorporating EMC products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Dell / EMC DD Series deduplication storage systems (including DD140, DD610, DD630, and DD670), which are appliances incorporating EMC's Data Domain Operating System,[23] including Global Compression™, and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

33. On information and belief, Dell and EMC have directly infringed and continue to infringe the '506 patent, for example, through their own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the

[21] https://www.emc.com/collateral/white-papers/h12145-intro-new-vnx-series-wp.pdf
[22] https://www.emc.com/collateral/white-papers/h11752-intro-to-XtremIO-array-wp.pdf
[23] http://www.emc.com/data-protection/data-domain/data-domain-operating-system.htm

'506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Dell and EMC use the Accused Instrumentality to practice infringing methods for their own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to Dell's and EMC's customers.

34. The Accused Instrumentality is a computer-implemented method for data compression. This system minimizes the amount of data transmitted over a network and stored on a backup device. The Accused Instrumentality employs several data compression techniques to achieve this goal.

35. The Accused Instrumentality analyzes data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types:[24]

[24] See https://community.emc.com/thread/203751

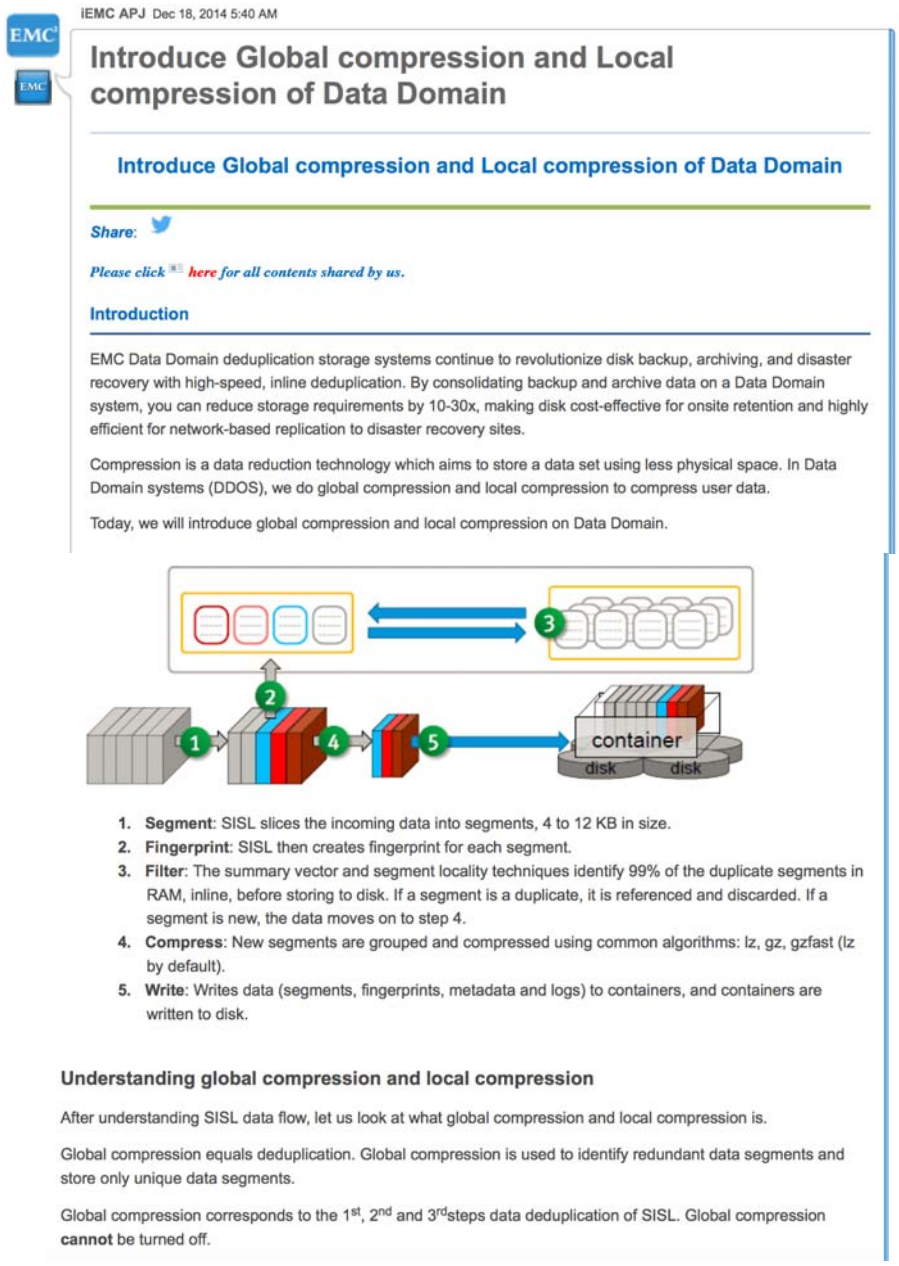

iEMC APJ Dec 18, 2014 5:40 AM

Introduce Global compression and Local compression of Data Domain

Introduce Global compression and Local compression of Data Domain

Share:

Please click here for all contents shared by us.

Introduction

EMC Data Domain deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. By consolidating backup and archive data on a Data Domain system, you can reduce storage requirements by 10-30x, making disk cost-effective for onsite retention and highly efficient for network-based replication to disaster recovery sites.

Compression is a data reduction technology which aims to store a data set using less physical space. In Data Domain systems (DDOS), we do global compression and local compression to compress user data.

Today, we will introduce global compression and local compression on Data Domain.

1. **Segment**: SISL slices the incoming data into segments, 4 to 12 KB in size.
2. **Fingerprint**: SISL then creates fingerprint for each segment.
3. **Filter**: The summary vector and segment locality techniques identify 99% of the duplicate segments in RAM, inline, before storing to disk. If a segment is a duplicate, it is referenced and discarded. If a segment is new, the data moves on to step 4.
4. **Compress**: New segments are grouped and compressed using common algorithms: lz, gz, gzfast (lz by default).
5. **Write**: Writes data (segments, fingerprints, metadata and logs) to containers, and containers are written to disk.

Understanding global compression and local compression

After understanding SISL data flow, let us look at what global compression and local compression is.

Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments.

Global compression corresponds to the 1st, 2nd and 3rdsteps data deduplication of SISL. Global compression **cannot** be turned off.

Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast).

Local compression corresponds to the 4th and 5th steps data deduplication of SISL. Local compression **can**be turned off.

The overall user data compression is the joint effort of global compression and local compression. This is the final result of deduplication. DDOS uses "compression ratio" to measure the effectiveness of its data compression.

36. The Accused Instrumentality performs content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified:[25]

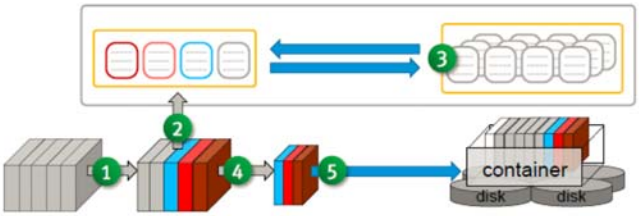

1. **Segment**: SISL slices the incoming data into segments, 4 to 12 KB in size.
2. **Fingerprint**: SISL then creates fingerprint for each segment.
3. **Filter**: The summary vector and segment locality techniques identify 99% of the duplicate segments in RAM, inline, before storing to disk. If a segment is a duplicate, it is referenced and discarded. If a segment is new, the data moves on to step 4.
4. **Compress**: New segments are grouped and compressed using common algorithms: lz, gz, gzfast (lz by default).
5. **Write**: Writes data (segments, fingerprints, metadata and logs) to containers, and containers are written to disk.

Understanding global compression and local compression

After understanding SISL data flow, let us look at what global compression and local compression is.

Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments.

Global compression corresponds to the 1st, 2nd and 3rdsteps data deduplication of SISL. Global compression **cannot** be turned off.

37. The Accused Instrumentality performs data compression with a single data compression encoder, if a data type of the data block is not identified:[26]

[25] See https://community.emc.com/thread/203751
[26] See https://community.emc.com/thread/203751

Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast).

Local compression corresponds to the 4th and 5th steps data deduplication of SISL. Local compression **can**be turned off.

The overall user data compression is the joint effort of global compression and local compression. This is the final result of deduplication. DDOS uses "compression ratio" to measure the effectiveness of its data compression.

38. In the Accused Instrumentality, analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block:[27]

[27] See https://community.emc.com/thread/203751

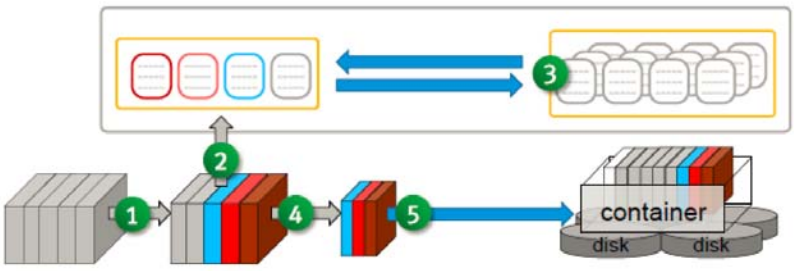

1. **Segment**: SISL slices the incoming data into segments, 4 to 12 KB in size.
2. **Fingerprint**: SISL then creates fingerprint for each segment.
3. **Filter**: The summary vector and segment locality techniques identify 99% of the duplicate segments in RAM, inline, before storing to disk. If a segment is a duplicate, it is referenced and discarded. If a segment is new, the data moves on to step 4.
4. **Compress**: New segments are grouped and compressed using common algorithms: lz, gz, gzfast (lz by default).
5. **Write**: Writes data (segments, fingerprints, metadata and logs) to containers, and containers are written to disk.

Understanding global compression and local compression

After understanding SISL data flow, let us look at what global compression and local compression is.

Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments.

Global compression corresponds to the 1st, 2nd and 3rdsteps data deduplication of SISL. Global compression **cannot** be turned off.

39. On information and belief, Dell and EMC also directly infringe and continue to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

40. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way. In particular, similar deduplication and compression technology used in EMC's Data Domain product is also used in Dell / EMC DD Series deduplication storage systems (including DD140, DD610, DD630, and DD670), which are appliances incorporating EMC's Data Domain Operating System,[28] including Global Compression™.

41. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

[28] http://www.emc.com/data-protection/data-domain/data-domain-operating-system.htm

42. On information and belief, Dell has had knowledge of the '506 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, Dell knew of the '506 patent and knew of its infringement, including by way of this lawsuit.

43. Upon information and belief, Dell's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Dell instructs customers of the Dell / EMC DD Series deduplication storage systems that "The Dell / EMC DD Series are mature backup to disk solutions with integrated deduplication. The solutions are designed to be easily incorporated into enterprise environments for customers who want to implement deduplication without changing their backup software. Data Domain technology has been built from the ground up to optimize Global Compression™ together with Stream Informed Segment Layout (SISL™) Scaling Architecture so that customers reap the benefits of both CPU performance scalability and reductions in backup media requirements."[29] For similar

[29] http://www.dell.com/downloads/global/products/pvaul/en/dell-emc-dd-series-brochure.pdf

reasons, Dell also induces its customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. Dell specifically intended and was aware that these normal and customary activities would infringe the '506 patent. Dell performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Dell engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Dell has induced and continue to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '506 patent, knowing that such use constitutes infringement of the '506 patent.

44. On information and belief, EMC has had knowledge of the '506 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, EMC knew of the '506 patent and knew of its infringement, including by way of this lawsuit.

45. Upon information and belief, EMC's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a

descriptor that is indicative of the data type of the data within the data block. For example, EMC instructs users of EMC Data Domain that "EMC Data Domain deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. … Compression is a data reduction technology which aims to store a data set using less physical space. In Data Domain systems (DDOS), we do global compression and local compression to compress user data. … Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments. … Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast)." [30] For similar reasons, EMC also induces its customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. EMC specifically intended and was aware that these normal and customary activities would infringe the '506 patent. EMC performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, EMC engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, EMC has induced and continue to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '506 patent, knowing that such use constitutes infringement of the '506 patent.

46. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Dell and EMC have injured Realtime and are liable to Realtime for infringement of the '506 patent pursuant to 35 U.S.C. § 271.

47. As a result Dell's and EMC's infringement of the '506 patent, Plaintiff

[30] See https://community.emc.com/thread/203751

Realtime is entitled to monetary damages in an amount adequate to compensate for Dell's and EMC's infringement, but in no event less than a reasonable royalty for the use made of the invention by Dell and EMC, together with interest and costs as fixed by the Court.

EMC Data Domain Boost (with Veeam Availability Suite™)

48. On information and belief, EMC and Veeam have made, used, offered for sale, sold and/or imported into the United States products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, EMC's and Veeam's compression products and services, such as, *e.g.*, EMC Data Domain Boost with Veeam Availability Suite™ v8,[31] and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

49. On information and belief, EMC and Veeam have directly infringed and continue to infringe the '506 patent, for example, through their own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the '506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, EMC and Veeam use the Accused Instrumentality to practice infringing methods for their own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to

[31] https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html

EMC's customers and Veeam's customers.

50. The Accused Instrumentality is a computer-implemented method for data compression. This system minimizes the amount of data transmitted over a network and stored on a backup device. The Accused Instrumentality employs several data compression techniques to achieve this goal. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent.").

51. The Accused Instrumentality analyzes data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for deduplication processing."); https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html:

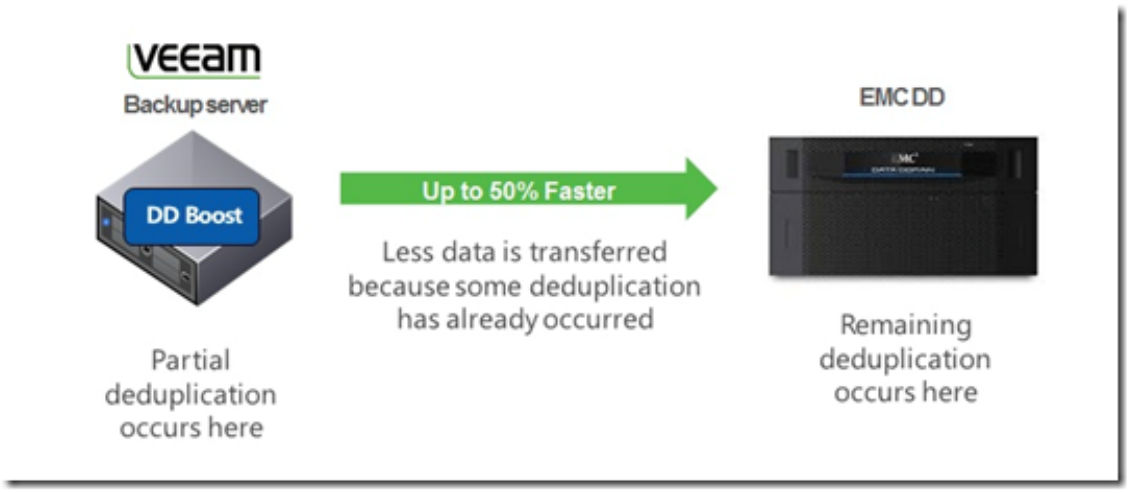

52. The Accused Instrumentality performs content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for deduplication processing. With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent."); https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html:

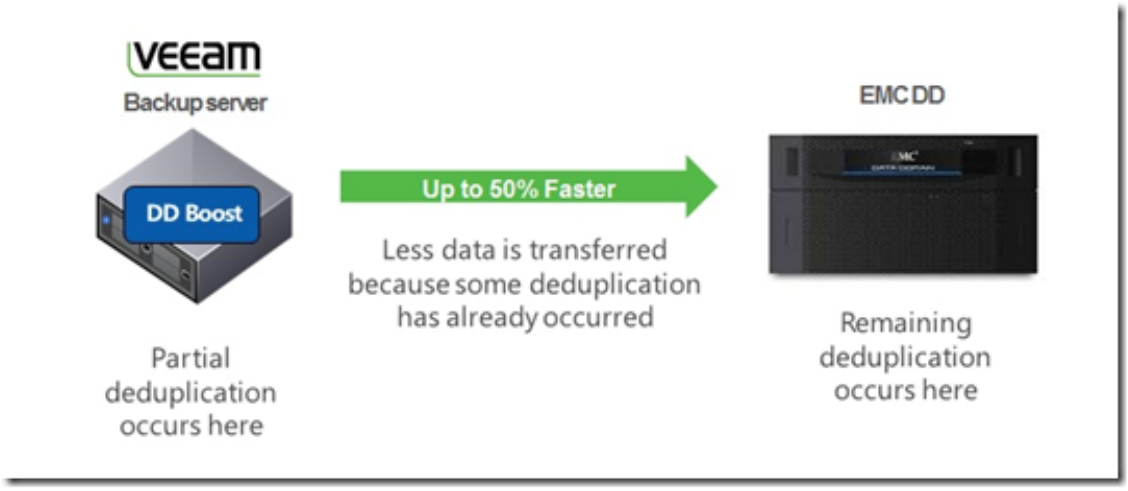

http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/k ("What is DD Boost? … Veeam will no longer send everything across the LAN to the DataDomain. Now, it is aware of what data is on the DataDomain, and takes part in the deduplication, and then will only send unique data.").

53. The Accused Instrumentality performs data compression with a single data compression encoder, if a data type of the data block is not identified. *See, e.g.*, http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/:

How to migrate:

This is actually a very straight forward process, as we don't edit the job, we create a new one.

- Setup DD Boost share on DataDomain
- Add DataDomain DD Boost enabled repository in Veeam
- Right click each job and choose to clone
- Edit newly cloned job, and point to DD Boost repository
- Edit job settings:
  - Enable inline deduplication
  - Compression level: Optimal
  - Storage Optimization: Local target

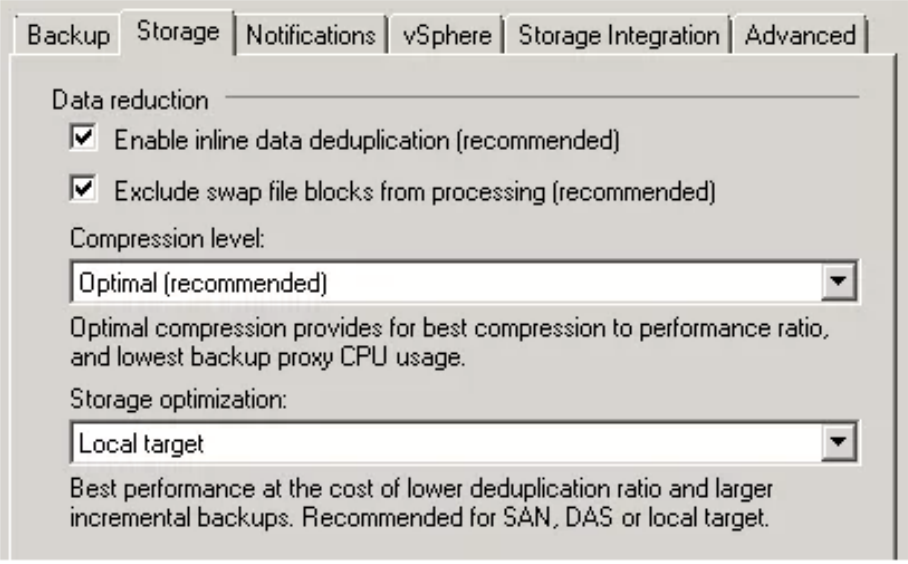

54. In the Accused Instrumentality analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for deduplication

processing. With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent."); https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html:

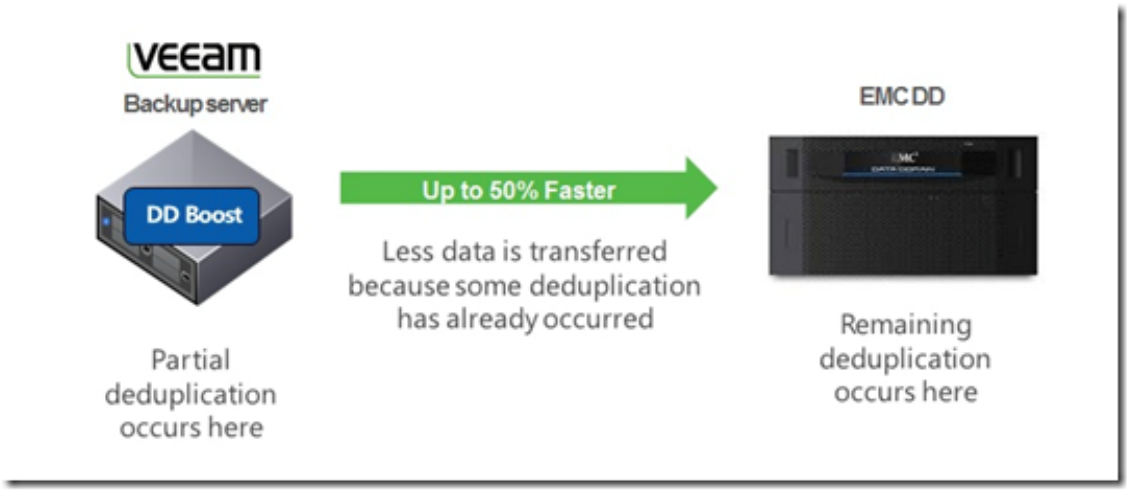

http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/k ("What is DD Boost? … Veeam will no longer send everything across the LAN to the DataDomain. Now, it is aware of what data is on the DataDomain, and takes part in the deduplication, and then will only send unique data.").

55. On information and belief, EMC and Veeam also directly infringe and continue to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

56. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

57. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

58. On information and belief, EMC and Veeam have had knowledge of the '506 patent since at least the filing of this Complaint or shortly thereafter, and on

information and belief, EMC and Veeam knew of the '506 patent and knew of their infringement, including by way of this lawsuit.

59. Upon information and belief, EMC and Veeam's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, EMC instructs users of EMC Data Domain Boost about the benefits of its deduplication processes. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for deduplication processing. With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent."). Veeam also instructs users of EMC Data Domain Boost with Veeam Availability Suite about the benefits of its deduplication processes. *See, e.g.*, https://www.veeam.com/blog/emc-data-domain-boost-coming-to-

veeam-availability-suite-v8.html:

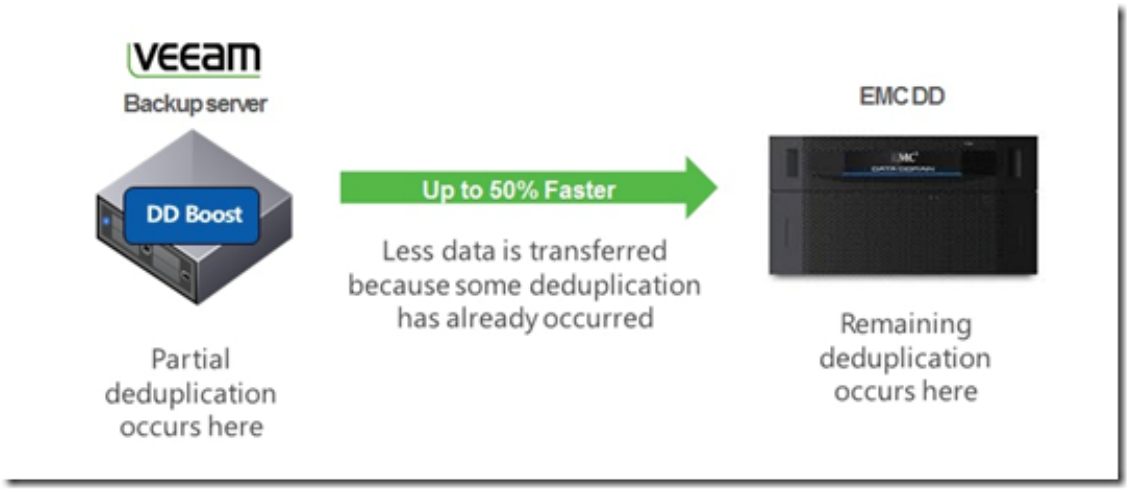

For similar reasons, EMC and Veeam also induce their customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. EMC and Veeam specifically intended and were aware that these normal and customary activities would infringe the '506 patent. EMC and Veeam performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, EMC and Veeam engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, EMC and Veeam have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '506 patent, knowing that such use constitutes infringement of the '506 patent.

60. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, EMC and Veeam have injured Realtime and are liable to Realtime for infringement of the '506 patent pursuant to 35 U.S.C. § 271.

61. As a result of EMC and Veeam's infringement of the '506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for EMC

and Veeam's infringement, but in no event less than a reasonable royalty for the use made of the invention by EMC and Veeam, together with interest and costs as fixed by the Court.

Veeam Backup & Replication

62. On information and belief, Veeam and Iland have made, used, offered for sale, sold and/or imported into the United States Veeam products that infringe the '506 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Veeam's compression products and services, such as, *e.g.*, Veeam Backup & Replication (which includes Veeam Cloud Connect), "iland Cloud Backup using Veeam Cloud Connect," and all versions and variations thereof since the issuance of the '506 patent ("Accused Instrumentality").

63. On information and belief, Veeam and Iland have directly infringed and continue to infringe the '506 patent, for example, through their own use and testing of the Accused Instrumentality to practice compression methods claimed by Claim 104 of the '506 patent, namely, a computer implemented method for compressing data, comprising: analyzing data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types; performing content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified; and performing data compression with a single data compression encoder, if a data type of the data block is not identified, wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. Upon information and belief, Veeam and Iland use the Accused Instrumentality to practice infringing methods for their own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to Veeam's customers and Iland's customers.

64. The Accused Instrumentality is a computer-implemented method for data compression. This system minimizes the amount of data transmitted over a network and stored on a backup device. The Accused Instrumentality employs several data compression techniques to achieve this goal.

65. The Accused Instrumentality analyzes data within a data block of an input data stream to identify one or more data types of the data block, the input data stream comprising a plurality of disparate data types. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html ("You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file.").

66. The Accused Instrumentality performs content dependent data compression with a content dependent data compression encoder if a data type of the data block is identified. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html ("You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file."); https://www.veeam.com/hyper-v-vmware-backup-deduplication-compression.html:

Fast, built-in deduplication

As its primary means of data reduction, Veeam uses a performance-optimized combination of deduplication and compression, allowing a single backup proxy to process data streams at 1GB/s during both VMware or Hyper-V backups and restores — up to10 times faster than the competition!

Veeam performs deduplication at the job level at both the source (i.e., at backup proxy) and the target (i.e., at backup repository):

- **Source-side deduplication** ensures that only unique data blocks not already present in the previous restore point are transferred across the network
- **Target-side deduplication** checks the received blocks against other virtual machine (VM) blocks already stored in the backup file, thus providing global deduplication across all VMs included in the backup job.

67. The Accused Instrumentality performs data compression with a single data compression encoder, if a data type of the data block is not identified. *See*, *e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html:

Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. Veeam Backup & Replication allows you to select one of the following compression levels:

- **None** compression level is recommended if you use storage devices with hardware compression and deduplication tools to store created backups.
- **Dedupe-friendly** is an optimized compression level for very low CPU usage. It is recommended if you want to decrease the proxy load.
- **Optimal** (default setting) is the recommended compression level providing the best ratio between the size of the backup file and time of the backup procedure.
- **High** compression level provides additional 10% compression ratio over **Optimal**, but at the cost of about 10x higher CPU usage.
- **Extreme** compression provides the smallest size of the backup file but reduces backup performance. We recommend that you run backup proxies on computers with modern multi-core CPUs (6 cores recommended) if you intend to use the extreme compression.

and https://www.veeam.com/hyper-v-vmware-backup-deduplication-compression.html:

Deduplication appliance-friendly compression — best of both worlds

Through its "dedupe-friendly" option, which uses a RLE (Run-Length Encoding) algorithm, Veeam Backup & Replication can send Veeam-deduplicated data to a backup appliance, and then have the backup appliance apply its own native deduplication to achieve even further data reduction. This allows you to have the best of both worlds — reducing bandwidth usage due to at-source data reduction, along with additional advanced deduplication by the storage.

68. In the Accused Instrumentality analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. *See* http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html ("You can apply deduplication when backing up multiple VMs that have similar data blocks (for

example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file. Depending on the type of storage you select as a backup target, Veeam Backup & Replication uses data blocks of different size to process VMs, which optimizes the size of a backup file and job performance.").

69. On information and belief, Veeam and Iland also directly infringe and continue to infringe other claims of the '506 patent, for similar reasons as explained above with respect to Claim 104 of the '506 patent.

70. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way. In particular, "iland Cloud Backup using Veeam Cloud Connect" [32] uses the same Veeam Cloud Connect that is a feature of Veeam Backup & Replication. [33]

71. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the methods claimed by the '506 patent.

72. On information and belief, Iland and Veeam have had knowledge of the '506 patent since at least the filing of this Complaint or shortly thereafter, and on information and belief, Iland and Veeam knew of the '506 patent and knew of their infringement, including by way of this lawsuit.

---

[32] http://www.iland.com/services/iland-cloud-backup-with-veeam/

[33] https://www.veeam.com/wp-cloud-connect-reference-architecture-veeam-backup-replication-v8.html ("One of the new and greatest features of Veeam® Backup & Replication™ v8 is Veeam Cloud Connect. With it, Veeam users can easily send backup copies offsite to remote locations managed by Veeam Service Providers. Also, Service Providers can use Cloud Connect to build their own remote repositories and offer their customers Backup Storage as a Service."); https://www.veeam.com/find-a-veeam-cloud-provider.html (identifying Iland as one of many "Veeam Cloud & Service Provider partners with datacenters in United States").

73. Upon information and belief, Iland and Veeam's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 104 of the '506 patent by practicing a computer implemented method comprising: receiving a data block in an uncompressed form, said data block being included in a data stream; analyzing data within the data block to determine a type of said data block; and compressing said data block to provide a compressed data block, wherein if one or more encoders are associated to said type, compressing said data block with at least one of said one or more encoders, otherwise compressing said data block with a default data compression encoder, and wherein the analyzing of the data within the data block to identify one or more data types excludes analyzing based only on a descriptor that is indicative of the data type of the data within the data block. For example, Veeam instructs users of Veeam Backup and Replication that, "To decrease traffic and disk space required for storing backup files, Veeam Backup & Replication provides mechanisms of data compression and deduplication. … Data compression decreases the size of created backups … Optimal (default setting) is the recommended compression level providing the best ratio between the size of the backup file and time of the backup procedure. … You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. … With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file." [34] For similar reasons, Veeam and Iland also induce their customers to use the Accused Instrumentalities to infringe other claims of the '506 patent. Veeam and Iland specifically intended and were aware that these normal and customary

[34] http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html

activities would infringe the '506 patent. Veeam and Iland performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '506 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Veeam and Iland engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Veeam and Iland have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '506 patent, knowing that such use constitutes infringement of the '506 patent.

74. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Veeam and Iland have injured Realtime and are liable to Realtime for infringement of the '506 patent pursuant to 35 U.S.C. § 271.

75. As a result Veeam and Iland's infringement of the '506 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Veeam and Iland's infringement, but in no event less than a reasonable royalty for the use made of the invention by Veeam and Iland, together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 9,054,728

76. Plaintiff Realtime realleges and incorporates by reference paragraphs 1-75 above, as if fully set forth herein.

77. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,054,728 ("the '728 Patent") entitled "Data compression systems and methods." The '728 Patent was duly and legally issued by the United States Patent and Trademark Office on June 9, 2015. A true and correct copy of the '728 Patent is included as Exhibit B.

Dell Rapid Recovery Software

78. On information and belief, Dell has, or will soon have, used, offered for sale, sold and/or imported into the United States Dell products that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Dell's compression products and services, such as, *e.g.*, the Rapid Recovery software product (v6.0.1),[35] the AppAssure v5.4 software product,[36] the Dell DL1000 Backup and Recovery Appliance,[37] the Dell PowerVault DL4000 Backup and Recovery Appliance,[38] the Dell DL4300 Backup and Recovery Appliance,[39] the Dell DR4100 Disk Backup Appliance, the Dell DR6000 Disk Backup Appliance,[40] the Dell DR2000v backup disk virtual appliance,[41] the Dell SonicWALL WAN Acceleration Virtual Appliance (WXA) 5000, and all versions and variations thereof since the issuance of the '728 patent ("Accused Instrumentality").

79. On information and belief, Dell has directly infringed and continues to infringe the '728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the '728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to

[35] http://software.dell.com/landing/6178
[36] http://documents.software.dell.com/appassure/5.4.1/user-guide/introduction-to-appassure-5/about-appassure-5
[37] http://software.dell.com/products/appassure-dl1000-backup-and-recovery-appliance/
[38] https://partnerdirect.dell.com/sites/channel/Documents/Dell-PowerVault-DL4000-Spec-Sheet.pdf
[39] http://software.dell.com/products/dl4300-backup-and-recovery-appliance/
[40] http://software.dell.com/products/dr-series-disk-backup-appliances/
[41] http://software.dell.com/products/dr2000v-virtual-backup-appliance/

perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, Dell uses the Accused Instrumentality, an infringing system, for its own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to Dell's customers.

80. The Accused Instrumentality is a system for compressing data, comprising a processor and one or more content dependent data compression encoders:

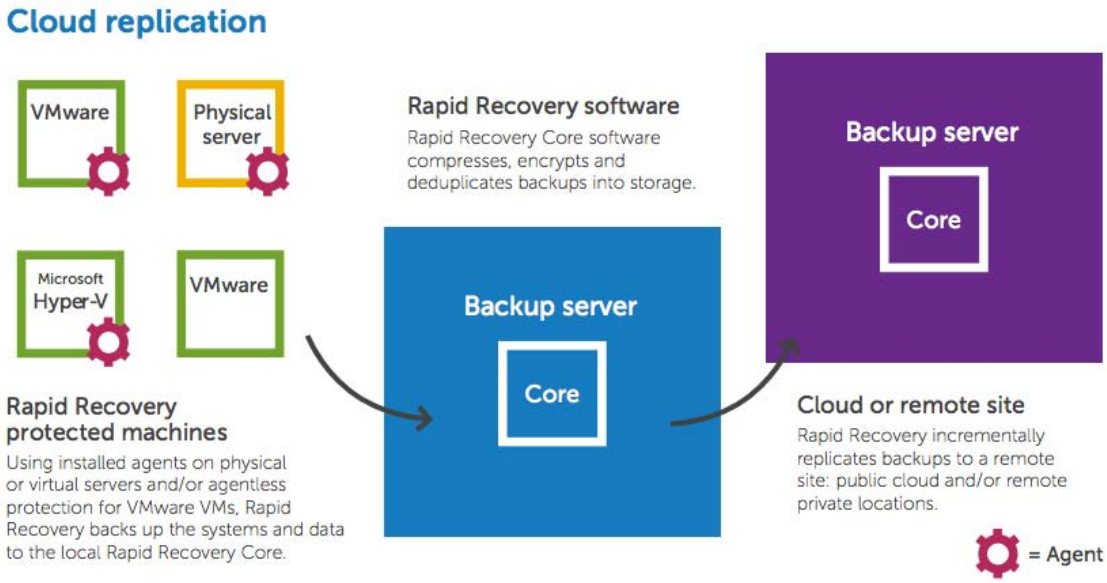

See http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf. See also http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf ("The new Rapid Recovery Repository (R3) is powered by the same leading-edge technology as Dell DR series backup and deduplication appliances."); http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf:

> The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

81. The Accused Instrumentality uses a single data compression encoder. *See, e.g.,* http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("Rapid Recovery feature list: … There is a new deduplication and compression engine from the DR target-based product set, meaning Dell now has a single dedupe engine across its portfolio."); http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf ("The Dell DR4000 has made further optimizations, and can actually dedupe and compress as part of the same inline process. This provides the benefits of compression without requiring that space be dedicated to staging uncompressed data.").

82. The Accused Instrumentality analyzes data within a data block to identify one or more parameter of the data, in this case, whether the data has been recognized as having been seen by the system before and where the analysis does not rely only on the descriptor. *See, e.g.,* http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf ("The R3, powered by the best-of-breed data deduplication technology platform used in Dell DR series appliances, can reside on separate Windows or Linux servers … Client-side data deduplication (also called source-side deduplication) uses the agent to remove redundant backup data before transmitting data to the R3. Using client-side data deduplication in tandem with target deduplication greatly reduces the amount of data sent over a LAN."); http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("AppAssure becomes the Rapid Recovery product under a Dell Data

Protection (DDP) brand. The product is block- and snapshot-based, only storing unique, changed blocks with a five-minute RPO."). See also http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf:

> The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

83. The Accused Instrumentality performs content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified. *See, e.g.*, http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf ("The R3, powered by the best-of-breed data deduplication technology platform used in Dell DR series appliances, can reside on separate Windows or Linux servers … Client-side data deduplication (also called source-side deduplication) uses the agent to remove redundant backup data before transmitting data to the R3. Using client-side data deduplication in tandem with target deduplication greatly reduces the amount of data sent over a LAN."); http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("AppAssure becomes the Rapid Recovery product under a Dell Data Protection (DDP) brand. The product is block- and snapshot-based, only storing unique, changed blocks with a five-minute RPO."). See also http://www.dell.com/downloads/global/products/pvaul/en/demystifying-deduplication.pdf:

The Dell DR4000 uses sophisticated pattern-matching algorithms to scan the next section of the data stream. This process determines the optimum chunk size for the incoming data and assigns the chunk a unique value, called a fingerprint. The Dell DR4000 uses industry-standard techniques to guarantee that a given fingerprint is assigned to single, unique, set of bits. The Dell DR4000 then checks the fingerprint against its dictionary. If the fingerprint corresponds to data already in the dictionary, the Dell DR4000 creates a pointer and adds one to the reference count. If the fingerprint does not correspond to data in the dictionary, the data is stored and the fingerprint is added to the dictionary. Having processed the chunk, the Dell DR4000 then continues processing the data stream.

84. The Accused Instrumentality performs data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified:

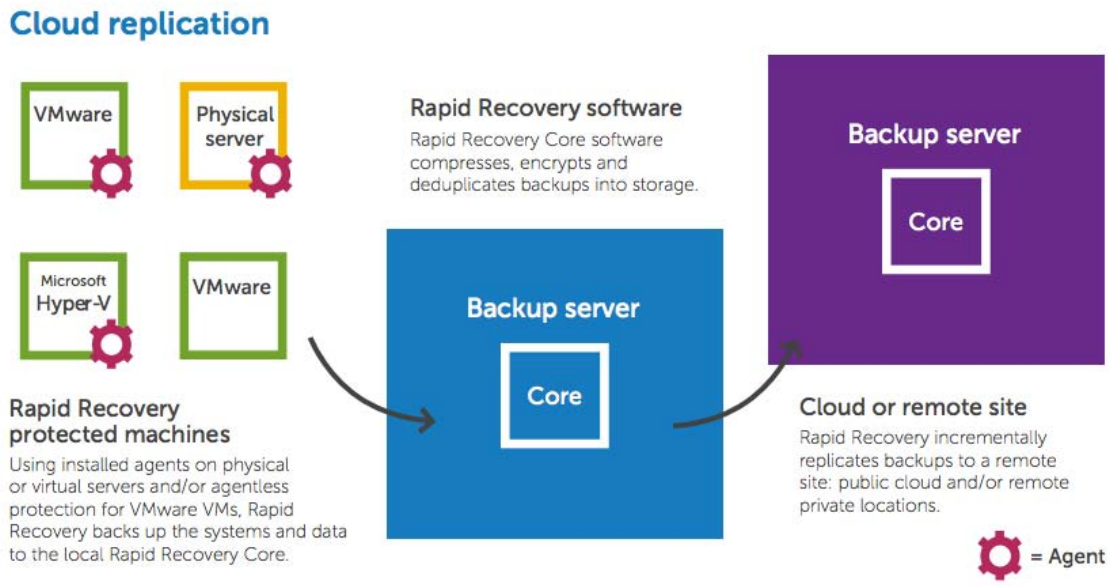

See http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf.

85. On information and belief, Dell also directly infringes and continues to infringe other claims of the '728 patent, for similar reasons as explained above with respect to Claim 1 of the '728 patent.

86. On information and belief, all of the Accused Instrumentalities operate in substantially the same way. In particular, similar deduplication and compression technology used in Dell's DR backup appliances (e.g. DR4100, DR6000, and DR2000v) is used in Dell's Rapid Recovery software, Dell's AppAssure software, and in Dell's

appliances built using Dell's AppAssure software (e.g. the Dell DL1000 Backup and Recovery Appliance, the Dell PowerVault DL4000 Backup and Recovery Appliance, and the Dell DL4300 Backup and Recovery Appliance). See, e.g., http://www.theregister.co.uk/2015/10/26/dell_begins_rebranding_appassure_becomes_rapid_recovery/ ("Rapid Recovery feature list: … There is a new deduplication and compression engine from the DR target-based product set, meaning Dell now has a single dedupe engine across its portfolio."); http://software.dell.com/products/appassure-dl1000-backup-and-recovery-appliance/ ("Built by Dell and powered by AppAssure, the DL1000 provides fast backup and restores on virtual machines, physical servers and in the cloud to enable local, offsite or disaster recovery."); http://software.dell.com/documents/dell-appassure-replication-technicalbrief-29890.pdf ("Dell™ AppAssure™ delivers advanced, flexible replication options to protect any organization, along with compression, deduplication and encryption"); http://documents.software.dell.com/appassure/5.4.3/user-guide-revision-b-english/configuring-the-appassure-core/managing-a-repository/modifying-repository-settings ("Enable Deduplication: Clear this checkbox to turn off deduplication, or select this checkbox to enable deduplication; Enable Compression: Clear this checkbox to turn off compression, or select this checkbox to enable compression.").

87. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '728 patent.

88. On information and belief, Dell has had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, Dell knew of the '728 patent and knew of its infringement, including by way of this lawsuit.

89. Upon information and belief, Dell's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and

technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, Dell instructs its customers that the "Rapid Recovery Core software compresses, encrypts and deduplicates backups into storage."[42] Dell specifically intended and was aware that these normal and customary activities would infringe the '728 patent. Dell performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Dell engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Dell has induced and continues to induce users of the accused products to use the Accused Instrumentalities in their ordinary and customary way to infringe the '728 patent, knowing that such use constitutes infringement of the '728 patent.

90. By making, using, offering for sale, selling and/or importing into the

[42] http://software.dell.com/docs/get-your-apps-back-in-business-technical-brief-102869.pdf

United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Dell has injured Realtime and is liable to Realtime for infringement of the '728 patent pursuant to 35 U.S.C. § 271.

91. As a result of Dell's infringement of the '728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Dell's infringement, but in no event less than a reasonable royalty for the use made of the invention by Dell, together with interest and costs as fixed by the Court.

EMC Data Domain

92. On information and belief, EMC has used, offered for sale, sold and/or imported into the United States EMC products that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, EMC's compression products and services, such as, *e.g.*, the EMC Data Domain product, the EMC VNX2 Series products (including the EMC VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 products), the EMC XtremIO Storage Array, and all versions and variations thereof since the issuance of the '728 patent ("Accused Instrumentality").

93. On information and belief, Dell has made, used, offered for sale, sold and/or imported into the United States combination products incorporating EMC products that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Dell / EMC DD Series deduplication storage systems (including DD140, DD610, DD630, and DD670), which are appliances incorporating EMC's Data Domain Operating System,[43] including Global Compression™, and all versions and variations thereof since the issuance of the '728 patent ("Accused Instrumentality").

94. On information and belief, Dell and EMC have directly infringed and

[43] http://www.emc.com/data-protection/data-domain/data-domain-operating-system.htm

continue to infringe the '728 patent, for example, through their own use and testing of the Accused Instrumentality, which constitutes systems for compressing data claimed by Claim 1 of the '728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, Dell and EMC use the Accused Instrumentality, an infringing system, for their own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to Dell's and EMC's customers.

95. The Accused Instrumentality is a system for compressing data, comprising a processor and one or more content dependent data compression encoders:

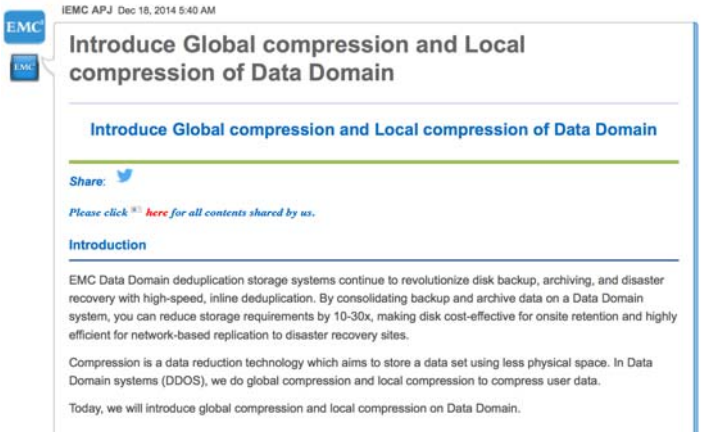

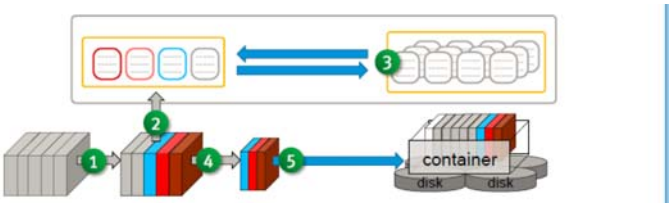

1. **Segment**: SISL slices the incoming data into segments, 4 to 12 KB in size.
2. **Fingerprint**: SISL then creates fingerprint for each segment.
3. **Filter**: The summary vector and segment locality techniques identify 99% of the duplicate segments in RAM, inline, before storing to disk. If a segment is a duplicate, it is referenced and discarded. If a segment is new, the data moves on to step 4.
4. **Compress**: New segments are grouped and compressed using common algorithms: lz, gz, gzfast (lz by default).
5. **Write**: Writes data (segments, fingerprints, metadata and logs) to containers, and containers are written to disk.

Understanding global compression and local compression

After understanding SISL data flow, let us look at what global compression and local compression is.

Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments.

Global compression corresponds to the 1st, 2nd and 3rdsteps data deduplication of SISL. Global compression **cannot** be turned off.

96. The Accused Instrumentality uses a single data compression encoder:[44]

Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast).

Local compression corresponds to the 4th and 5th steps data deduplication of SISL. Local compression **can**be turned off.

The overall user data compression is the joint effort of global compression and local compression. This is the final result of deduplication. DDOS uses "compression ratio" to measure the effectiveness of its data compression.

97. The Accused Instrumentality analyzes data within a data block to identify one or more parameter of the data, in this case, whether the data has been recognized as having been seen by the system before and where the analysis does not rely only on the descriptor: [45]

[44] https://community.emc.com/thread/203751
[45] https://community.emc.com/thread/203751

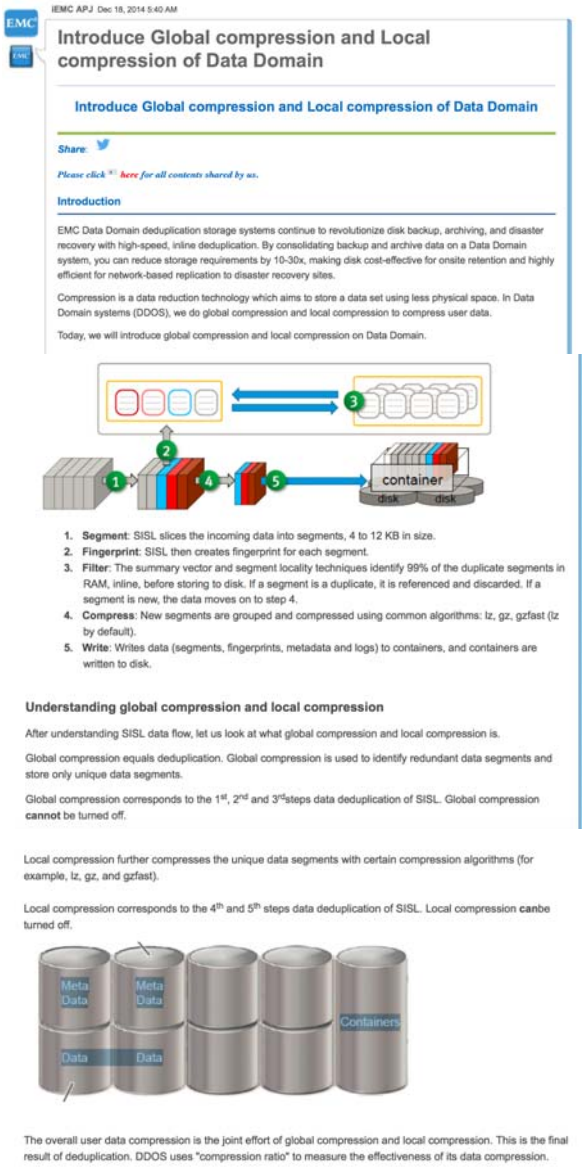

IEMC APJ Dec 18, 2014 5:40 AM

Introduce Global compression and Local compression of Data Domain

Introduce Global compression and Local compression of Data Domain

Share:

Please click [*] here for all contents shared by us.

Introduction

EMC Data Domain deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. By consolidating backup and archive data on a Data Domain system, you can reduce storage requirements by 10-30x, making disk cost-effective for onsite retention and highly efficient for network-based replication to disaster recovery sites.

Compression is a data reduction technology which aims to store a data set using less physical space. In Data Domain systems (DDOS), we do global compression and local compression to compress user data.

Today, we will introduce global compression and local compression on Data Domain.

1. **Segment**: SISL slices the incoming data into segments, 4 to 12 KB in size.
2. **Fingerprint**: SISL then creates fingerprint for each segment.
3. **Filter**: The summary vector and segment locality techniques identify 99% of the duplicate segments in RAM, inline, before storing to disk. If a segment is a duplicate, it is referenced and discarded. If a segment is new, the data moves on to step 4.
4. **Compress**: New segments are grouped and compressed using common algorithms: lz, gz, gzfast (lz by default).
5. **Write**: Writes data (segments, fingerprints, metadata and logs) to containers, and containers are written to disk.

Understanding global compression and local compression

After understanding SISL data flow, let us look at what global compression and local compression is.

Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments.

Global compression corresponds to the 1st, 2nd and 3rdsteps data deduplication of SISL. Global compression **cannot** be turned off.

Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast).

Local compression corresponds to the 4th and 5th steps data deduplication of SISL. Local compression **can**be turned off.

The overall user data compression is the joint effort of global compression and local compression. This is the final result of deduplication. DDOS uses "compression ratio" to measure the effectiveness of its data compression.

98. The Accused Instrumentality performs content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified: [46]

[46] https://community.emc.com/thread/203751

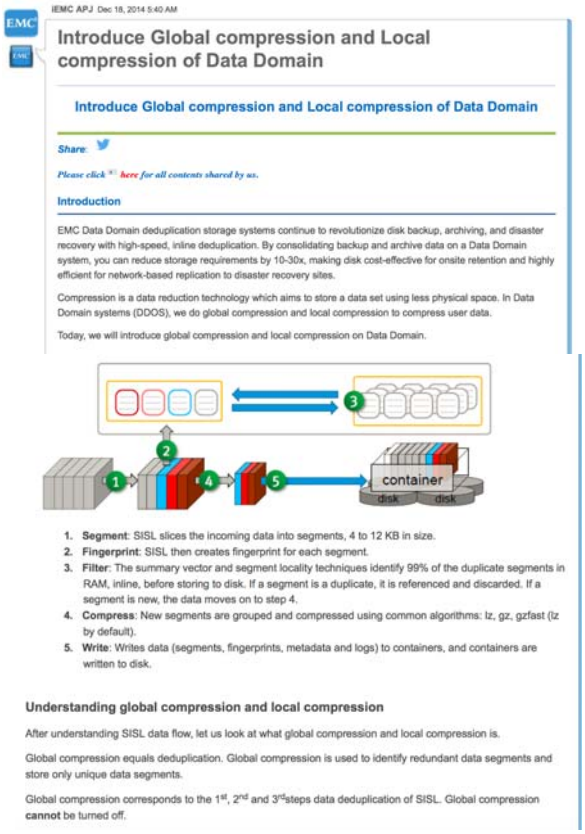

IEMC APJ Dec 18, 2014 5:40 AM

Introduce Global compression and Local compression of Data Domain

Introduce Global compression and Local compression of Data Domain

Share:

Please click [*] here for all contents shared by us.

Introduction

EMC Data Domain deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. By consolidating backup and archive data on a Data Domain system, you can reduce storage requirements by 10-30x, making disk cost-effective for onsite retention and highly efficient for network-based replication to disaster recovery sites.

Compression is a data reduction technology which aims to store a data set using less physical space. In Data Domain systems (DDOS), we do global compression and local compression to compress user data.

Today, we will introduce global compression and local compression on Data Domain.

1. **Segment**: SISL slices the incoming data into segments, 4 to 12 KB in size.
2. **Fingerprint**: SISL then creates fingerprint for each segment.
3. **Filter**: The summary vector and segment locality techniques identify 99% of the duplicate segments in RAM, inline, before storing to disk. If a segment is a duplicate, it is referenced and discarded. If a segment is new, the data moves on to step 4.
4. **Compress**: New segments are grouped and compressed using common algorithms: lz, gz, gzfast (lz by default).
5. **Write**: Writes data (segments, fingerprints, metadata and logs) to containers, and containers are written to disk.

Understanding global compression and local compression

After understanding SISL data flow, let us look at what global compression and local compression is.

Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments.

Global compression corresponds to the 1st, 2nd and 3rdsteps data deduplication of SISL. Global compression **cannot** be turned off.

99. The Accused Instrumentality performs data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified: [47]

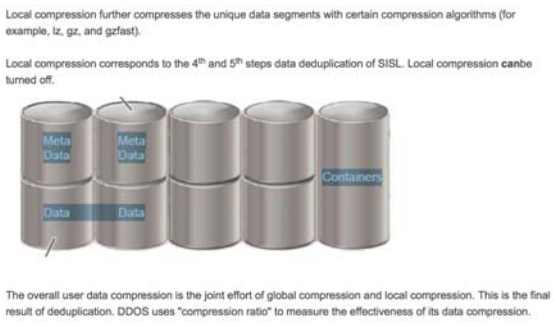

Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast).

Local compression corresponds to the 4th and 5th steps data deduplication of SISL. Local compression **can**be turned off.

The overall user data compression is the joint effort of global compression and local compression. This is the final result of deduplication. DDOS uses "compression ratio" to measure the effectiveness of its data compression.

100. On information and belief, Dell and EMC also directly infringe and continue to infringe other claims of the '728 patent, for similar reasons as explained above with respect to Claim 1 of the '728 patent.

101. On information and belief, all of the Accused Instrumentalities infringe

[47] https://community.emc.com/thread/203751

the '728 patent in substantially the same way. In particular, similar deduplication and compression technology used in EMC's Data Domain product is also used in Dell / EMC DD Series deduplication storage systems (including DD140, DD610, DD630, and DD670), which are appliances incorporating EMC's Data Domain Operating System,[48] including Global Compression™.

102. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the '728 patent.

103. On information and belief, Dell has had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, Dell knew of the '728 patent and knew of its infringement, including by way of this lawsuit.

104. Upon information and belief, Dell's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using systems for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if

[48] http://www.emc.com/data-protection/data-domain/data-domain-operating-system.htm

the one or more parameters or attributes of the data are not identified. For example, Dell instructs customers of the Dell / EMC DD Series deduplication storage systems that "The Dell / EMC DD Series are mature backup to disk solutions with integrated deduplication. The solutions are designed to be easily incorporated into enterprise environments for customers who want to implement deduplication without changing their backup software. Data Domain technology has been built from the ground up to optimize Global Compression™ together with Stream Informed Segment Layout (SISL™) Scaling Architecture so that customers reap the benefits of both CPU performance scalability and reductions in backup media requirements."[49] For similar reasons, Dell also induces its customers to use the Accused Instrumentalities to infringe other claims of the '728 patent. Dell specifically intended and was aware that these normal and customary activities would infringe the '728 patent. Dell performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Dell engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Dell has induced and continue to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '728 patent, knowing that such use constitutes infringement of the '728 patent.

105. On information and belief, EMC has had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, EMC knew of the '728 patent and knew of its infringement, including by way of this lawsuit.

106. Upon information and belief, EMC's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and

[49] http://www.dell.com/downloads/global/products/pvaul/en/dell-emc-dd-series-brochure.pdf

technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using systems for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, EMC instructs users of EMC Data Domain that "EMC Data Domain deduplication storage systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication. … Compression is a data reduction technology which aims to store a data set using less physical space. In Data Domain systems (DDOS), we do global compression and local compression to compress user data. … Global compression equals deduplication. Global compression is used to identify redundant data segments and store only unique data segments. … Local compression further compresses the unique data segments with certain compression algorithms (for example, lz, gz, and gzfast)." [50] For similar reasons, EMC also induces its customers to use the Accused Instrumentalities to infringe other claims of the '728 patent. EMC specifically intended and was aware that these normal and customary activities would infringe the '728 patent. EMC performed the acts that constitute induced infringement, and would induce actual

[50] See https://community.emc.com/thread/203751

infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, EMC engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, EMC has induced and continue to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the '728 patent, knowing that such use constitutes infringement of the '728 patent.

107. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Dell and EMC have injured Realtime and are liable to Realtime for infringement of the '728 patent pursuant to 35 U.S.C. § 271.

108. As a result Dell's and EMC's infringement of the '728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Dell's and EMC's infringement, but in no event less than a reasonable royalty for the use made of the invention by Dell and EMC, together with interest and costs as fixed by the Court.

EMC Data Domain Boost (with Veeam Availability Suite™)

109. On information and belief, EMC and Veeam have made, used, offered for sale, sold and/or imported into the United States products that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, EMC's and Veeam's compression products and services, such as, *e.g.*, EMC Data Domain Boost with Veeam Availability Suite™ v8,[51] and all versions and variations thereof since the issuance of the '728 patent ("Accused Instrumentality").

110. On information and belief, EMC and Veeam have directly infringed and continue to infringe the '728 patent, for example, through their own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by

[51] https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html

Claim 1 of the '728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, EMC and Veeam have used the Accused Instrumentality, an infringing system, for their own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to EMC's customers and Dell's customers.

111. The Accused Instrumentality is a system for compressing data, comprising a processor and one or more content dependent data compression encoders. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent.").

112. The Accused Instrumentality uses a single data compression encoder. *See, e.g.*, http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/:

How to migrate:

This is actually a very straight forward process, as we don't edit the job, we create a new one.

- Setup DD Boost share on DataDomain
- Add DataDomain DD Boost enabled repository in Veeam
- Right click each job and choose to clone
- Edit newly cloned job, and point to DD Boost repository
- Edit job settings:
  - Enable inline deduplication
  - Compression level: Optimal
  - Storage Optimization: Local target

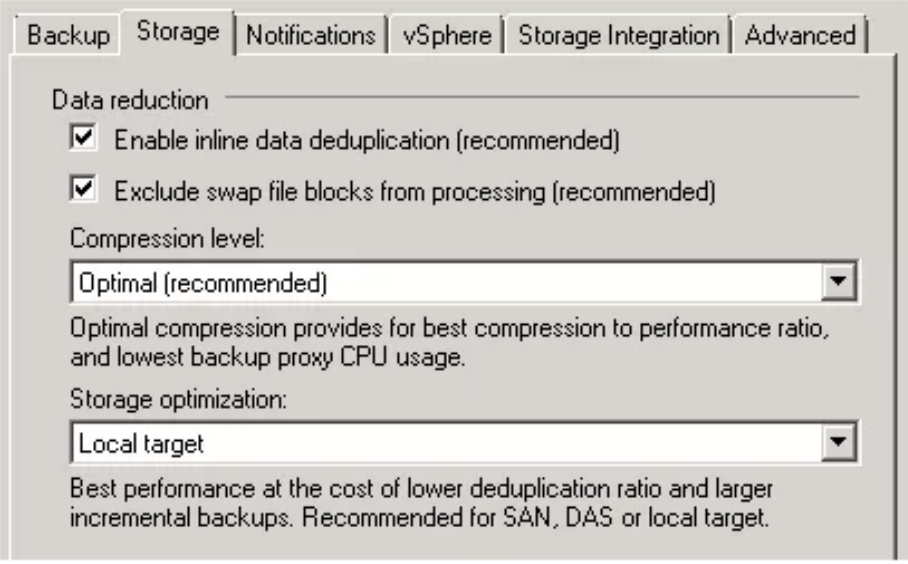

113. The Accused Instrumentality analyzes data within a data block to identify one or more parameter of the data, in this case, whether the data has been recognized as having been seen by the system before and where the analysis does not rely only on the descriptor. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for

deduplication processing. With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent."); https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html:

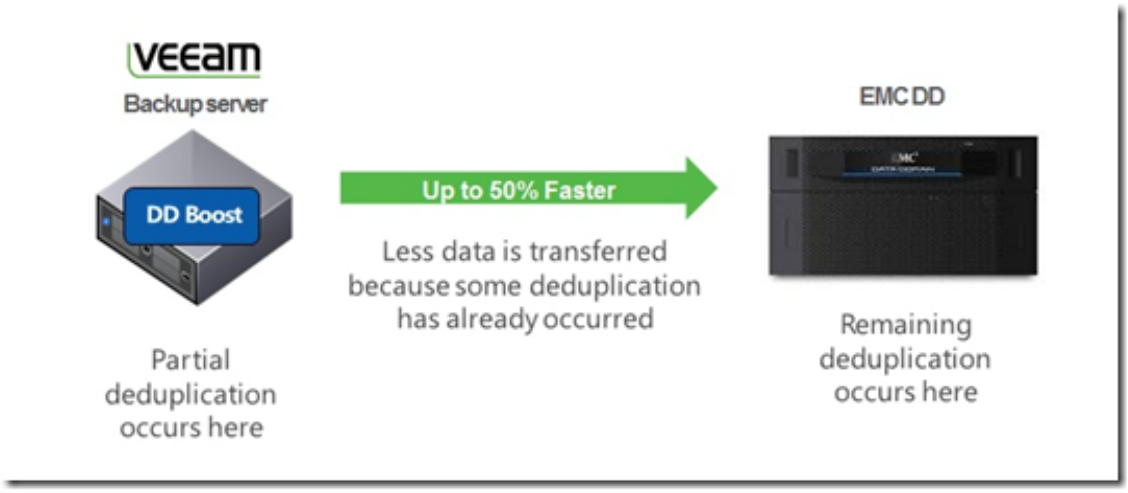

http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/k ("What is DD Boost? … Veeam will no longer send everything across the LAN to the DataDomain. Now, it is aware of what data is on the DataDomain, and takes part in the deduplication, and then will only send unique data.").

114. The Accused Instrumentality performs content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for deduplication processing. With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically increases the aggregate throughput by 50 percent

and reduces the amount of data transferred over the network by 80 to 99 percent."); https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html:

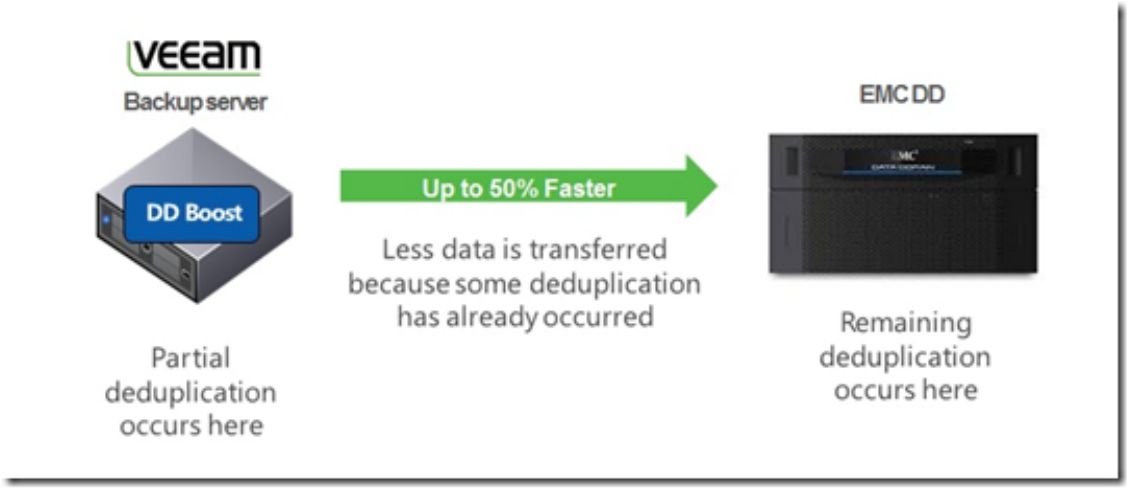

http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/k ("What is DD Boost? … Veeam will no longer send everything across the LAN to the DataDomain. Now, it is aware of what data is on the DataDomain, and takes part in the deduplication, and then will only send unique data.").

115. The Accused Instrumentality performs data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. *See, e.g.*, http://tsmith.co/2015/veeam-and-datadomain-using-dd-boost/:

How to migrate:

This is actually a very straight forward process, as we don't edit the job, we create a new one.

- Setup DD Boost share on DataDomain
- Add DataDomain DD Boost enabled repository in Veeam
- Right click each job and choose to clone
- Edit newly cloned job, and point to DD Boost repository
- Edit job settings:
  - Enable inline deduplication
  - Compression level: Optimal
  - Storage Optimization: Local target

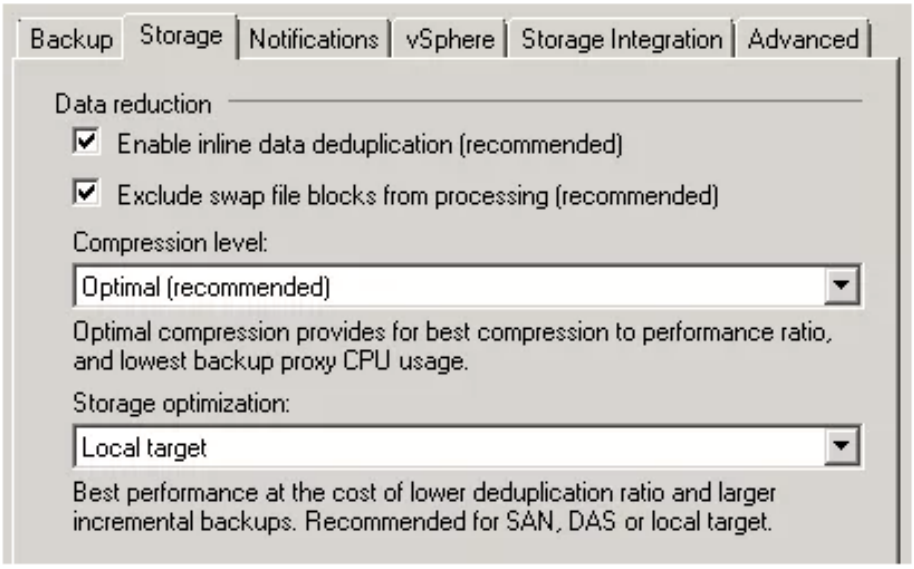

116. On information and belief, EMC and Veeam also directly infringe and continue to infringe other claims of the '728 patent, for similar reasons as explained above with respect to Claim 1 of the '728 patent.

117. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way.

118. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '728 patent.

119. On information and belief, EMC and Veeam have had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, EMC and Veeam knew of the '728 patent and knew of their infringement, including by way of this lawsuit.

120. Upon information and belief, EMC and Veeam's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, EMC instructs users of EMC Data Domain Boost about the benefits of its deduplication processes. *See, e.g.*, http://www.emc.com/data-protection/data-domain/data-domain-boost.htm#!details ("EMC Data Domain Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery. Without it, the backup server or application client would send all data—unique or redundant—to an EMC Data Domain system for deduplication processing. With Data Domain Boost, the server only sends unique data segments to a Data Domain system. This dramatically

increases the aggregate throughput by 50 percent and reduces the amount of data transferred over the network by 80 to 99 percent."). Veeam also instructs users of EMC Data Domain Boost with Veeam Availability Suite about the benefits of its deduplication processes. *See, e.g.*, https://www.veeam.com/blog/emc-data-domain-boost-coming-to-veeam-availability-suite-v8.html:

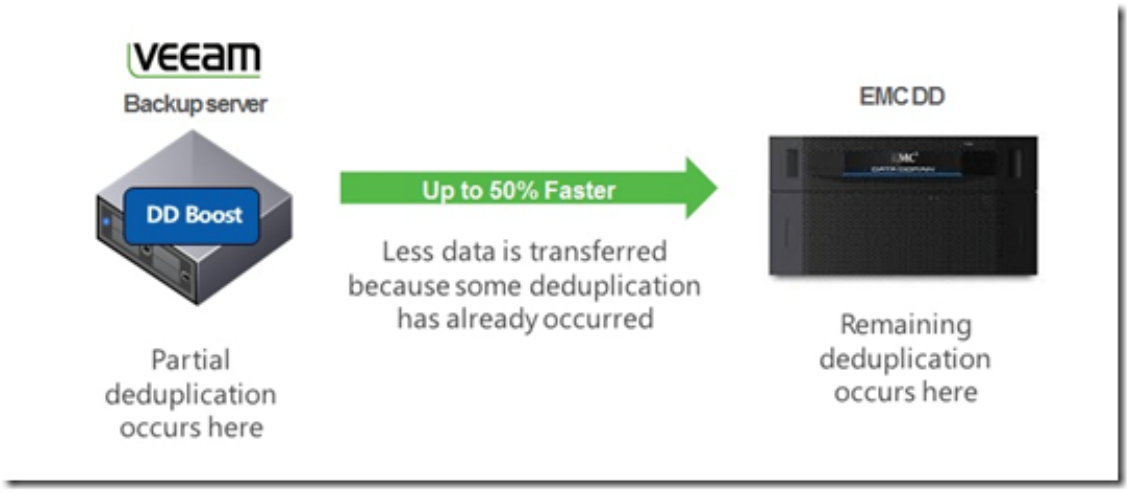

For similar reasons, EMC and Veeam also induce their customers to use the Accused Instrumentalities to infringe other claims of the '728 patent. EMC and Veeam specifically intended and were aware that these normal and customary activities would infringe the '728 patent. EMC and Veeam performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, EMC and Veeam engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, EMC and Veeam have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '728 patent, knowing that such use constitutes infringement of the '728 patent.

121. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the

Accused Instrumentalities' compression features, EMC and Veeam have injured Realtime and are liable to Realtime for infringement of the '728 patent pursuant to 35 U.S.C. § 271.

122. As a result of EMC and Veeam's infringement of the '728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for EMC and Veeam's infringement, but in no event less than a reasonable royalty for the use made of the invention by EMC and Veeam, together with interest and costs as fixed by the Court.

Veeam Backup & Replication

123. On information and belief, Veeam and Iland have made, used, offered for sale, sold and/or imported into the United States Veeam products that infringe the '728 patent, and continues to do so. By way of illustrative example, these infringing products include, without limitation, Veeam's compression products and services, such as, *e.g.*, Veeam Backup & Replication (which includes Veeam Cloud Connect), "iland Cloud Backup using Veeam Cloud Connect," and all versions and variations thereof since the issuance of the '728 patent ("Accused Instrumentality").

124. On information and belief, Veeam and Iland have directly infringed and continue to infringe the '728 patent, for example, through its own use and testing of the Accused Instrumentality, which constitute systems for compressing data claimed by Claim 1 of the '728 patent, comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data

compression encoder, if the one or more parameters or attributes of the data are not identified. Upon information and belief, Veeam and Iland use the Accused Instrumentality, an infringing system, for their own internal non-testing business purposes, while testing the Accused Instrumentality, and while providing technical support and repair services for the Accused Instrumentality to Veeam's customers and Iland's customers.

125. The Accused Instrumentality is a system for compressing data, comprising a processor and one or more content dependent data compression encoders. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html ("You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file.").

126. The Accused Instrumentality uses a single data compression encoder. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html:

Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. Veeam Backup & Replication allows you to select one of the following compression levels:

- **None** compression level is recommended if you use storage devices with hardware compression and deduplication tools to store created backups.
- **Dedupe-friendly** is an optimized compression level for very low CPU usage. It is recommended if you want to decrease the proxy load.
- **Optimal** (default setting) is the recommended compression level providing the best ratio between the size of the backup file and time of the backup procedure.
- **High** compression level provides additional 10% compression ratio over **Optimal**, but at the cost of about 10x higher CPU usage.
- **Extreme** compression provides the smallest size of the backup file but reduces backup performance. We recommend that you run backup proxies on computers with modern multi-core CPUs (6 cores recommended) if you intend to use the extreme compression.

and https://www.veeam.com/hyper-v-vmware-backup-deduplication-compression.html:

Deduplication appliance-friendly compression — best of both worlds

Through its "dedupe-friendly" option, which uses a RLE (Run-Length Encoding) algorithm, Veeam Backup & Replication can send Veeam-deduplicated data to a backup appliance, and then have the backup appliance apply its own native deduplication to achieve even further data reduction. This allows you to have the best of both worlds — reducing bandwidth usage due to at-source data reduction, along with additional advanced deduplication by the storage.

127. The Accused Instrumentality analyzes data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html ("You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file."); https://www.veeam.com/hyper-v-vmware-backup-deduplication-compression.html:

Fast, built-in deduplication

As its primary means of data reduction, Veeam uses a performance-optimized combination of deduplication and compression, allowing a single backup proxy to process data streams at 1GB/s during both VMware or Hyper-V backups and restores — up to10 times faster than the competition!

Veeam performs deduplication at the job level at both the source (i.e., at backup proxy) and the target (i.e., at backup repository):

- **Source-side deduplication** ensures that only unique data blocks not already present in the previous restore point are transferred across the network
- **Target-side deduplication** checks the received blocks against other virtual machine (VM) blocks already stored in the backup file, thus providing global deduplication across all VMs included in the backup job.

128. The Accused Instrumentality performs content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html ("You can apply deduplication when backing up multiple VMs that have similar data blocks (for

example, if VMs were created from the same template) or great amount of free space on their logical disks. Veeam Backup & Replication does not store zero byte blocks or space that has been pre-allocated but not used. With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file."); https://www.veeam.com/hyper-v-vmware-backup-deduplication-compression.html:

Fast, built-in deduplication

As its primary means of data reduction, Veeam uses a performance-optimized combination of deduplication and compression, allowing a single backup proxy to process data streams at 1GB/s during both VMware or Hyper-V backups and restores — up to10 times faster than the competition!

Veeam performs deduplication at the job level at both the source (i.e., at backup proxy) and the target (i.e., at backup repository):

- **Source-side deduplication** ensures that only unique data blocks not already present in the previous restore point are transferred across the network
- **Target-side deduplication** checks the received blocks against other virtual machine (VM) blocks already stored in the backup file, thus providing global deduplication across all VMs included in the backup job.

129. The Accused Instrumentality performs data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. *See, e.g.*, http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html:

Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. Veeam Backup & Replication allows you to select one of the following compression levels:

- **None** compression level is recommended if you use storage devices with hardware compression and deduplication tools to store created backups.
- **Dedupe-friendly** is an optimized compression level for very low CPU usage. It is recommended if you want to decrease the proxy load.
- **Optimal** (default setting) is the recommended compression level providing the best ratio between the size of the backup file and time of the backup procedure.
- **High** compression level provides additional 10% compression ratio over **Optimal**, but at the cost of about 10x higher CPU usage.
- **Extreme** compression provides the smallest size of the backup file but reduces backup performance. We recommend that you run backup proxies on computers with modern multi-core CPUs (6 cores recommended) if you intend to use the extreme compression.

and https://www.veeam.com/hyper-v-vmware-backup-deduplication-compression.html:

Deduplication appliance-friendly compression — best of both worlds

Through its "dedupe-friendly" option, which uses a RLE (Run-Length Encoding) algorithm, Veeam Backup & Replication can send Veeam-deduplicated data to a backup appliance, and then have the backup appliance apply its own native deduplication to achieve even further data reduction. This allows you to have the best of both worlds — reducing bandwidth usage due to at-source data reduction, along with additional advanced deduplication by the storage.

130. On information and belief, Iland and Veeam also directly infringe and continue to infringe other claims of the '728 patent, for similar reasons as explained above with respect to Claim 1 of the '728 patent.

131. On information and belief, all of the Accused Instrumentalities perform the claimed methods in substantially the same way. In particular, "iland Cloud Backup using Veeam Cloud Connect" [52] uses the same Veeam Cloud Connect that is a feature of Veeam Backup & Replication. [53]

132. On information and belief, use of the Accused Instrumentality in its ordinary and customary fashion results in infringement of the systems claimed by the '728 patent.

133. On information and belief, Iland and Veeam have had knowledge of the '728 patent since at least the filing of the original Complaint or shortly thereafter, and on information and belief, Iland and Veeam knew of the '728 patent and knew of their infringement, including by way of this lawsuit.

134. Upon information and belief, Iland and Veeam's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced

[52] http://www.iland.com/services/iland-cloud-backup-with-veeam/

[53] https://www.veeam.com/wp-cloud-connect-reference-architecture-veeam-backup-replication-v8.html ("One of the new and greatest features of Veeam® Backup & Replication™ v8 is Veeam Cloud Connect. With it, Veeam users can easily send backup copies offsite to remote locations managed by Veeam Service Providers. Also, Service Providers can use Cloud Connect to build their own remote repositories and offer their customers Backup Storage as a Service."); https://www.veeam.com/find-a-veeam-cloud-provider.html (identifying Iland as one of many "Veeam Cloud & Service Provider partners with datacenters in United States").

and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe the '728 patent by making or using a system for compressing data comprising a processor; one or more content dependent data compression encoders; and a single data compression encoder; wherein the processor is configured: to analyze data within a data block to identify one or more parameters or attributes of the data wherein the analyzing of the data within the data block to identify the one or more parameters or attributes of the data excludes analyzing based solely on a descriptor that is indicative of the one or more parameters or attributes of the data within the data block; to perform content dependent data compression with the one or more content dependent data compression encoders if the one or more parameters or attributes of the data are identified; and to perform data compression with the single data compression encoder, if the one or more parameters or attributes of the data are not identified. For example, Veeam instructs users of Veeam Backup and Replication that, "To decrease traffic and disk space required for storing backup files, Veeam Backup & Replication provides mechanisms of data compression and deduplication. … Data compression decreases the size of created backups … Optimal (default setting) is the recommended compression level providing the best ratio between the size of the backup file and time of the backup procedure. … You can apply deduplication when backing up multiple VMs that have similar data blocks (for example, if VMs were created from the same template) or great amount of free space on their logical disks. … With deduplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file." [54] For similar reasons, Veeam and Iland also induce their customers to use the Accused Instrumentalities to infringe other claims of the '728 patent. Veeam and Iland specifically intended and were aware that these normal and customary activities would infringe the '728 patent. Veeam and Iland performed the acts that

[54] http://helpcenter.veeam.com/backup/80/hyperv/compression_deduplication.html

constitute induced infringement, and would induce actual infringement, with the knowledge of the '728 patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, Veeam and Iland engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, Veeam and Iland have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their ordinary and customary way to infringe the '728 patent, knowing that such use constitutes infringement of the '728 patent.

135. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, Iland and Veeam have injured Realtime and are liable to Realtime for infringement of the '728 patent pursuant to 35 U.S.C. § 271.

136. As a result of Iland and Veeam's infringement of the '728 patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for Iland and Veeam's infringement, but in no event less than a reasonable royalty for the use made of the invention by Iland and Veeam, together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff Realtime respectfully requests that this Court enter:

a. A judgment in favor of Plaintiff that Dell, EMC, Iland, and Veeam have infringed, either literally and/or under the doctrine of equivalents, the '506 patent and the '728 patent;

b. A judgment and order requiring Dell, EMC, Iland, and Veeam to pay Plaintiff its damages, costs, expenses, and prejudgment and post-judgment interest for their infringement of the '506 patent and the '728 patent, as provided under 35 U.S.C. § 284;

c. A judgment and order requiring Dell, EMC, Iland, and Veeam to provide an accounting and to pay supplemental damages to Realtime, including without limitation, prejudgment and post-judgment interest;

d. A judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees against Dell, EMC, Iland, and Veeam; and

e. Any and all other relief as the Court may deem appropriate and just under the circumstances.

DEMAND FOR JURY TRIAL

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

Dated: February 26, 2016

Respectfully submitted,

/s/ Marc A. Fenster by permission Claire Abernathy Henry

Marc A. Fenster (CA SBN 181067)
Reza Mirzaie (CA SBN 246953)
Brian D. Ledahl (CA SBN 186579)
Jeffrey Z.Y. Liao (CA SBN 288994)
C. Jay Chung (CA SBN 252794)
RUSS AUGUST & KABAT
12424 Wilshire Boulevard, 12th Floor
Los Angeles, CA 90025
(310) 826-7474
mfenster@raklaw.com
rmirzaie@raklaw.com
bledahl@raklaw.com
jliao@raklaw.com
jchung@raklaw.com

T. John Ward, Jr.
Texas State Bar No. 00794818
E-mail: jw@wsfirm.com
Claire Abernathy Henry
Texas State Bar No. 24053063

E-mail: claire@wsfirm.com
WARD, SMITH & HILL, PLLC
1127 Judson Road, Ste 220
Longview, Texas 75601
(903) 757-6400 (telephone)
(903) 757-2323 (facsimile)

Attorneys for Plaintiff
Realtime Data LLC d/b/a IXO